

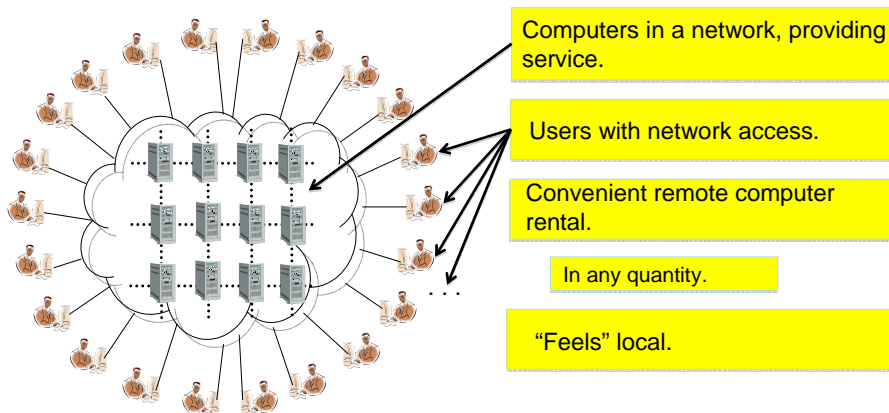
# Cloud Computing: Some Implications for Key Management

June 8, 2009

Lee Badger

*For those viewing via webcast, please submit questions for this presentation to [kmwquestions@nist.gov](mailto:kmwquestions@nist.gov)*




## Cloud Computing: Still Being Defined



A technical or business innovation?

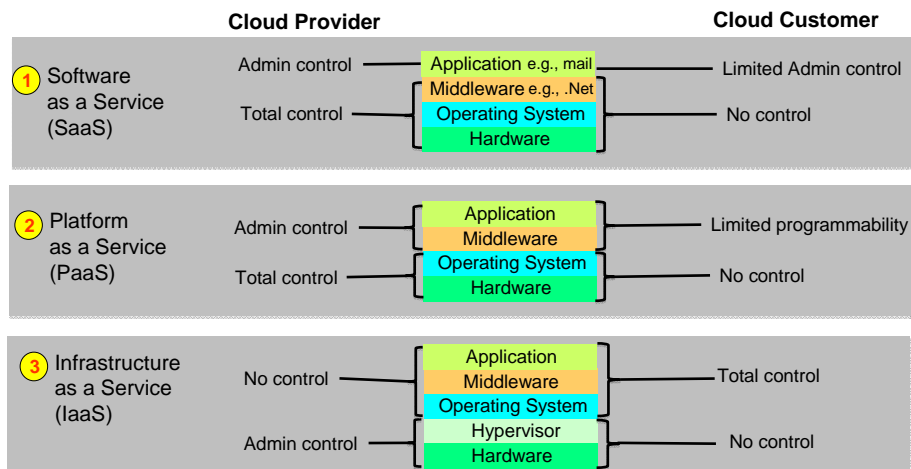
# NIST Working Cloud Definition (1 of 3)

## 5 Key Characteristics



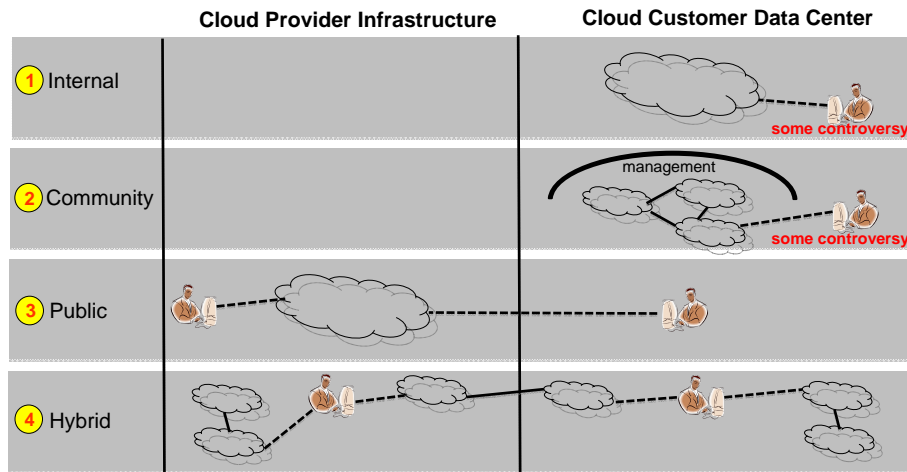
# NIST Working Cloud Definition (2 of 3)

## 3 Deployment Models



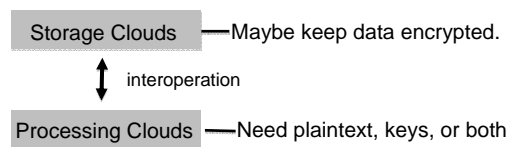
# NIST Working Cloud Definition (3 of 3)

## 4 Delivery Models



## A few more Cloud Aspects

- 1 There are two basic kinds of clouds:



- 2 Both **require** extremely **fast & reliable & secure & low-cost networking**.

- 3 Clouds are a good fit for very large scale processing/storage, using new algorithms:


Map Reduce  
Bigtable  
Google File System

## Why Care?

- ① Clouds capture economies of scale. In 2006, large data center cost advantage: 5.7 to 7.1 times.  
Credit: [Ham], summarized in [Arm].

- ② Clouds can locate computing close to low-cost power, land.

- ③ Clouds make it easier to provision for actual (fluctuating) loads.



Provide the illusion of infinite resources (to customers). If there is demand diversity.

- ④ Clouds enable rapid, low commitment, infrastructure for short term (or new) projects.

## Who Cares?

- Amazon
  - Microsoft
  - Google App Engine
  - Salesforce
  - IBM blue cloud
  - Vmware
  - Sun/Oracle
  - Force.com
  - ...
  - Open Cloud Consortium
  - Vivek Kundra, fed CIO
  - Distributed Management Task Force
- And more.**

Credit: <http://www.johnmwillis.com/cloud-computing/cloud-vendors-a-to-z-revised/>

## Sample Cloud Interface (Amazon EC2)

**43 customer-callable functions:**  
(SOAP, Query, command-line)

Operate your own cloud resources.  
All operations digitally signed.

Amazon Machine Image: 4

Key Pairs: 3

Availability Zones/Regions: 2

Monitoring: 2

Images: 2

Block Store Management: 8  
(snapshot)

Reserved Instances: 3

IP Addresses management: 5

Security Groups: 5

Instances: 4

MS Windows specific: 3

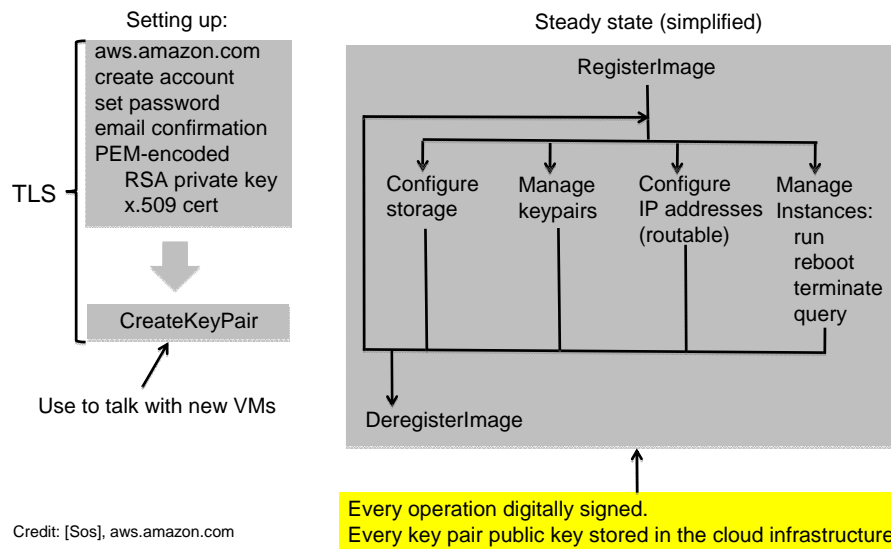
Credit: aws.amazon.com

(not an endorsement)

Information Technology Laboratory  
Computer Security Division

**NIST**  
National Institute of  
Standards and Technology

## A Quick Trip Through the (simplified) API



Credit: [Sos], aws.amazon.com

Information Technology Laboratory  
Computer Security Division

**NIST**  
National Institute of  
Standards and Technology

# Key Management

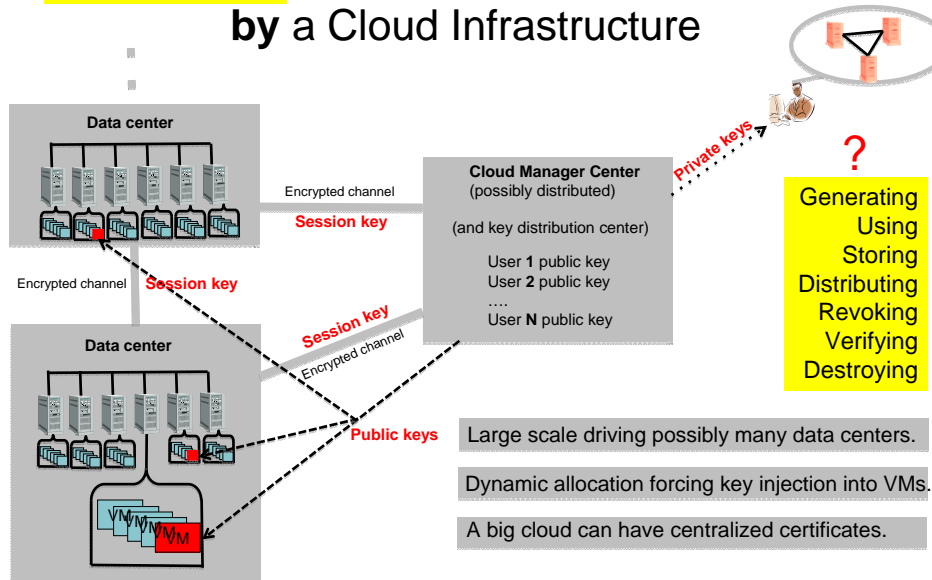
Key Management:

Generating keys	Using keys	Storing keys
Distributing keys	Revoking keys	...
Verifying keys	Destroying keys	

**Bold assertion:** there are two basic scenarios:

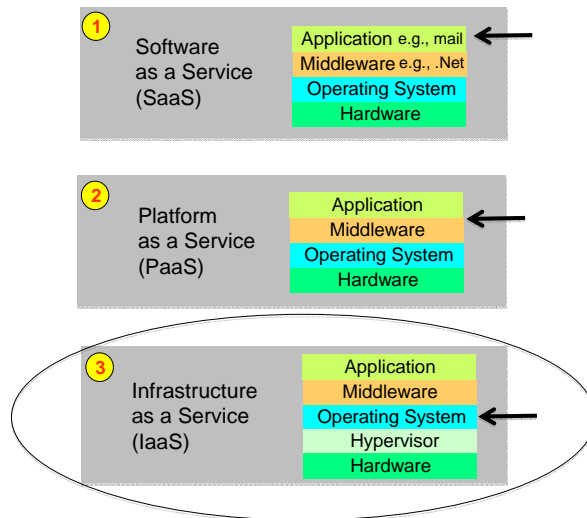
- 1 Key management conducted by the cloud infrastructure itself.
- 2 Key management conducted by computations that have been entrusted to run in the cloud infrastructure

## Speculation: How Keys Might Get Managed by a Cloud Infrastructure




Credit: [Nur] for inspiration via the Eucalyptus system, and aws.amazon.com.

## Execution Environments for Cloud-Hosted Key Management



## Hardware Virtualization



- A simple picture!
- But implementation is complex.
- Virtual Machines (VMs) can be:
  - suspended/copied/moved/lost/recovered.


# Hardware Virtualization

## (Simple View)

## Terminology

Guest OS : runs only on VMM  
Host OS : runs only on HW  
Domain : virtual machine on VMM  
Hypervisor : virtual machine monitor

**Issue:** Deprivileging



In 2000, the Pentium had 250 instructions, 18 that were not virtualizable. Binary translation, however, could virtualization them anyway.

## Making x86 Virtualizable Using Extra Hardware

# Intel 64

Intel version of x86-64

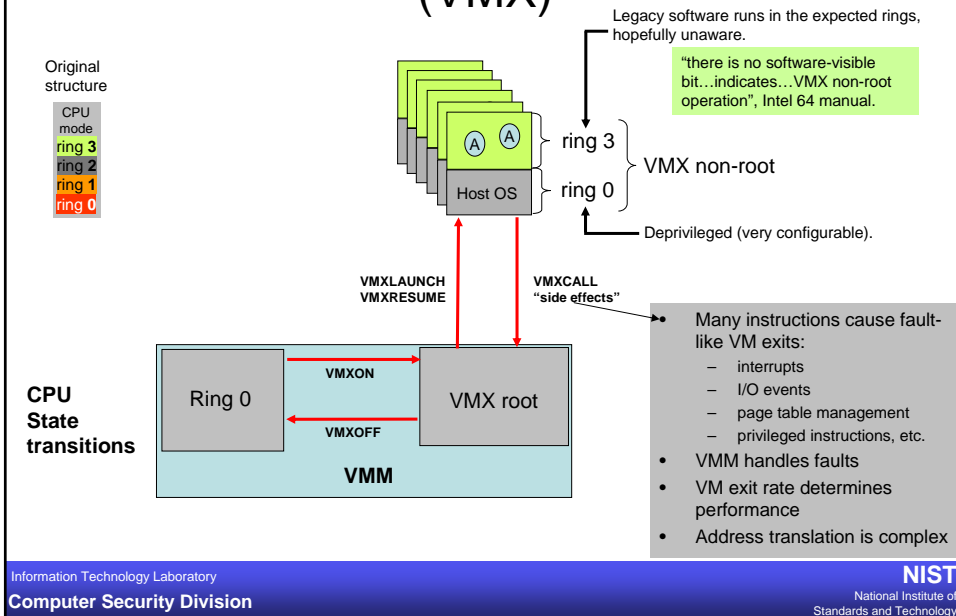
contains ~595 instructions.

Hardware extensions  
make the instruction set  
virtualizable


<b>Floating Point</b>					
Data	17				
Arithmetic	26				
Compare	14				
Transcendental	8				
Constants	7				
Control	20				
State management	2				
	<b>94</b>				
<b>SIMD</b>					
MMX	47				
SSE	62				
SSE2	69				
SSE3	13				
SSSE3	32				
SSE4	54				
	<b>277</b>				
<b>General Purpose</b>					
Data transfer	32				
Arithmetic	18				
Logical	4				
Shift/rotate	9				
Bit/byte	23				
Control transfer	31				
String	18				
I/O	8				
Enter/leave	2				
Flag control	11				
Segment register	5				
Misc	6				
	<b>167</b>				
<b>VT-x Extensions</b>		<b>12</b>			
<b>Safe mode</b>		<b>1</b>			
<b>Hardware extensions make the instruction set virtualizable</b>			<b>System</b>	<b>34</b>	
			<b>64-bit mode</b>	<b>10</b>	



# Intel Virtual Machine Extensions (VMX)



## How Complex is Virtualization?



# VMM Implementation Quality Should Not be Assumed

In 2007, Tavis Ormandy subjected 6 virtualization systems to guided random testing of their invalid instruction handling and I/O emulation.

Bochs

178k SLOC

QEMU

373k SLOC

VMWare

Xen

910k SLOC

Anonymous 1

Anonymous 2

All of the systems failed the tests, most with “arbitrary execution” failures.

Device emulation was a particular area of vulnerability.

For details, see: [taviso.decsystem.org/virtsec.pdf](http://taviso.decsystem.org/virtsec.pdf)

Reference: “An Empirical Study into the Security Exposures to Host of Hostile Virtualized Environments,”  
by Travis Ormandy, [taviso.decsystem.org/virtsec.pdf](http://taviso.decsystem.org/virtsec.pdf)  
Code counts generated using David A. Wheeler’s “SLOCCount” tool.

## Some Cloud Implications

For the cloud infrastructure itself:

Cloud infrastructures can centralize certificate hierarchies, at scale.

Time-based customer eviction may assist with key revocation, destruction.

Within a cloud: one scheme, one owner, one codebase.

Clouds can manage/control (e.g., not lose) VMs.

For computations run in the cloud:

Keys need a safe harbor in the cloud.

Trusted Platform Module (TPM) hard to virtualize.

Remote attestation may not work.

But users may be able to leverage the cloud infrastructure as a trusted party.

E.g., to rely on VM sanitization if promised.

# References

- [Arm] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. Above the Clouds: A Berkeley View of Cloud Computing. [www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html](http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html)
- [aws] Amazon Web Services, [aws.amazon.com](http://aws.amazon.com).
- [Ham] J. Hamilton. Internet-Scale Service Efficiency. N Large-Scale Distributed Systems and Middleware Workshop (Sept. 2008). See also: <http://perspectives.mvdirona.com/2008/09/16/InternetScaleServiceEfficiency.aspx>
- [Sos] Robert Sosinski. Starting Amazon EC2 with Mac OS X. <http://www.robertsosinski.com/2008/01/26/starting-amazon-ec2-with-mac-os-x/>
- [Nur] D. Nurmi, R. Wolski, C. Grzegorzcyk, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov. The Eucalyptus Open-source Cloud-computing System. [open.eucalyptus.com/documents/nurmi\\_et\\_al\\_eucalyptus\\_open\\_source\\_cloud\\_computing\\_system-cca\\_2008.pdf](http://open.eucalyptus.com/documents/nurmi_et_al_eucalyptus_open_source_cloud_computing_system-cca_2008.pdf)
- [Wil] <http://www.johnmwillis.com/cloud-computing/cloud-vendors-a-to-z-revised/>
- [Sch] B. Schneier. Applied Cryptography. ISBN 0-471-59756-2. 1993.
- <http://www.dmtf.org/about/cloud-incubator>
- [http://commons.wikimedia.org/wiki/Main\\_Page](http://commons.wikimedia.org/wiki/Main_Page). Images.
- [Orm] Tavis Ormandy. "An Empirical Study into the Security Exposures to Host of Hostile Virtualized Environments," [tavis.decsystem.org/virtsec.pdf](http://tavis.decsystem.org/virtsec.pdf)
- [Gon] Gonzalez-Barahona et al. "Counting Potatos: The Size of Debian 2.2". [People.debian.org/~jgb/debian-counting](http://people.debian.org/~jgb/debian-counting).
- [Whe] David Wheeler. More than a Gigabuck: Estimating GNU/Linux's Size. <http://www.dwheeler.com/sloc/redhat71-v1/redhat71sloc.html>
- Intel 64 reference manual. <http://www.intel.com/products/processor/manuals/>